Hardened Networks: Incremental Upgrading of the Internet for Attack Resilience

Shu Zhang
Arizona State University
Tempe, AZ 85287
shu.zhang@asu.edu

Partha Dasgupta
Arizona State University
Tempe, AZ 85287
partha@asu.edu

Abstract

Network security is conventionally implemented at the edge of the network (such as SSL, SSh), or router-based filtering. They require the awareness from the users and the understanding of the complicated configuration. They do little to provide resilience to network attacks.

In this paper, we discuss a different approach to enhance of Network Security. We use smarter routers to build security mechanisms (source authentication, flow control, encryption) into the fabric of the network. It allows for incremental upgrading as well as compatibility with all current protocols. Since the security mechanisms are at the router level, there is no impact on the end user. We also show some implementation and simulation results.

1. INTRODUCTION

When the current de-facto standard TCP/IP protocol suite was designed in 1979, public access to the network layer was not part of the design. Hence, the security features were not of major concern. As a result, the suite has plenty of vulnerabilities [1] — making the Internet a playground for hackers.

Even though TCP/IP forms the backbone of today's Internet, it lacks the most basic mechanisms for security, such as authentication or encryption. IP blindly delivers any packet to any specified IP. It is also easy to set fake control information in IP headers or sniff without being detected.

The current state of the art in network security uses techniques that work on the "edge of the network." Many security protocols are being used to secure the Internet communications, such as SSL [6], SSh [7], and IPSec [5].

These protocols require hosts to be explicitly configured (VPN [5] is transparent to the end applications, but the client computer needs to be configured). In addition, these techniques do little to provide resilience to a plethora of attacks on TCP/IP that are designed to exploit their vulnerabilities. The Attack prevention techniques rely heavily on firewalls [4]. The use of firewalls to prevent spoofing and scanning has had some success, but it failed to protect against more sophisticated attacks (DDoS [10]), and also restricts the network activities of the internal users.

The goal of our research is to make the Internet secure and resilient to attacks. *Our solution is counter to the conventional approach.*. Instead of running protocols on the hosts or filtering programs[3] on firewalls, our solution builds the hardening at the router level using cryptographic schemes and additional control points (for enhanced flow control), that is, there is *no* change to the protocols and protocol stacks at the hosts. Thus, the applications *do not* have to be changed. All TCP/IP enabled devices benefit from the hardened network. Even the network itself is better protected against the DoS attacks due to the source authentication and flow control.

The hardening of the network infrastructure is attained by replacing regular network routers with "*hardened routers*" that can perform a few extra functions, in addition to routing. *Completely revamping the Internet network protocol is obviously infeasible*; hence in our approach, the hardened routers are completely compatible with the existing routers and can be *incrementally* introduced – making deployment feasible. We call this approach "Network Hardening."

2. OUR APPROACH

2.1. OVERVIEW OF HARDENED NETWORK

As stated earlier, we harden the Internet by replacing routers with hardened routers. For simplicity, we assume that the Internet is a set of Autonomous Systems (ASes).

The key idea of "hardening" is the enhancement of some functions in a router. This enhancement is cryptography-based, including both encryption and authentication. The encryption of IP payloads secures the privacy of the Internet communications. Cryptographic techniques permit the construction of a trusted subsystem with distributed reach, which is crucial for identifying the attackers. .

We classify the routers in an AS into three categories:

Border routers: These are on the inter-AS path, and run Border Gateway Protocol for inter-AS routing.

Access routers: The routers are directly connected to the end host machines inside the same AS.

Core Routers: All routers, other than Border routers and Access routers, are Core Routers.

To create a Hardened AS, we harden the *Border Routers* and *Access Routers*, so that these routers can do encryption and information gathering. Core Routers are left untouched.

A hardened router is a regular router, which has the following additional abilities: sign/verify and encrypt/decrypt packets when needed; provide traffic statistics; filter/restrict packets when instructed by AS controller.

To facilitate global flow monitoring, router configuration management and digital certificate issuing, we add a "Hardened AS Controller" to each hardened AS.

2.2. BENEFITS

To illustrate the Hardened Network, consider the network shown in Fig. 2.

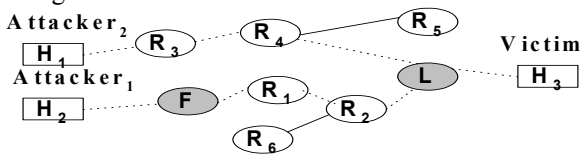

Figure. 2. Connections in Hardened Network

This is a hybrid network consisting of normal routers (R_1, R_2 …) and two hardened routers F and L. H_1 and H_2 are hosts involved in an attack and H_3 is the victim. F is the first hardened router between H_2 and H_3, which means there is no hardened router between H_2 and F. L is the last Hardened Router meaning no Hardened Router between L and H_3.

When F receives an IP packet, it marks the packet with its IP address and encrypts the payload and the first byte of the IP address using the key shared with L before it forwards it. In order to enhance the performance, we only require all the first hardened routers to mark and encrypt each IP packet. All the hardened routers between the first hardened router and the last hardened router simply forward all the packets. So when the marked and encrypted packet gets to L, L decrypts the payload using the shared key and verifies the IP address. In this way, we can guarantee that the identification in the marker field is not repudiated by F.

Consider the effect of hardening when H_1 and H_2 send out the DDoS attack traffic to H_3, indicated as dotted lines. The flow control component in L detects the congestion and packet dropping. According to the identification information, L immediately traces the attack traffic from H_2 back to F and from H_1 back to itself. Then, it informs F about the attack. After that, both L and F start dropping the attack traffic.

The attack traffic from H_2 can be stopped closer to the source, but not that from H_1. When more routers are hardened, we can stop the attack more effectively. Even in this network, we can still reduce the load to victim, so that it can serve more good traffic. Hence, only the two hardened routers present can ensure that the traffic flowing through them is resilient to attacks. Due to the signature, the command from L to drop the traffic cannot be forged or be denied.

Another advantage of this system is the fact that the detection and traceback are not done at the victim, but close to the victim, from where most of the attack traffic passes. The traceback result cannot be denial. While most the current DDoS traceback techniques ([11][12][13]), the traceback is done at the victim network when it is under extreme stress. Also they need some form of authentication to verify the traceback information.

Since the traffic flow between the first and the last hardened routers is encrypted, the packets are immune to sniffing without installing theVPN software.

When to encrypt/decrypt and how to perform key exchange and the secure communications among Hardened ASes are the key issues of "Hardened Network". The scope of this paper is limited to address the solutions to these problems, but not to defeat attacks, which was illustrated in [9].

3.HARDENED NETWORK

The hardened network is composed of Hardened ASes. Fig. 3 shows the components of each hardened AS.

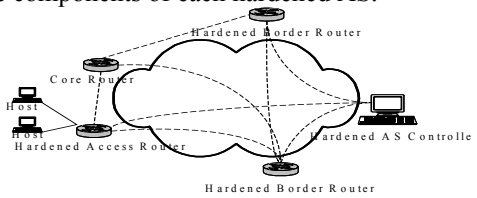

Figure 3. Hardened AS

3.1. THE HARDENED ROUTERS

There are two types of hardened routers: the "Hardened Border routers" and "Hardened Access routers."

Hardened Border Router: The hardened border router consists of a Local-Key-Base-In, a Local-Key-Base-Out, a Manager Process, a Key Exchange Agent, and a Local-HRIB, shown in Fig 4. The Local-Key-Base-In/out stores all the decryption/encryption keys. The Manager process manages both Key Bases. The Key Exchange Agent is responsible for exchanging the encryption key with Hardened AS Controllers. The Local-HRIB is used to save all the routing information and other information related to the hardened network that have been acquired by exchanging routing information.

Hardened Access Router: The hardened access router handles traffic originating from, or destined to, the hosts directly connected to them. The Local-Key-Base-In for the encryption keys stores the keys shared by the two local access routers. Apart from those components present in a hardened border router, the hardened access router has one more database, the Local-Info-Base, which stores all the host information of the local hardened AS.

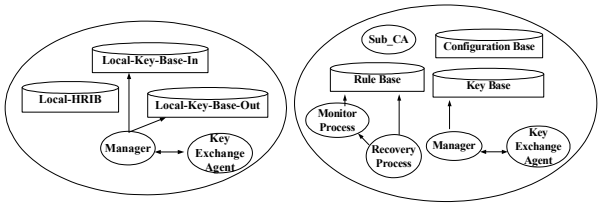

Figure 4. Hardened Border Router (L) Hardened AS Controller(R)

3.2. HARDENED AS CONTROLLER

The hardened AS Controller is a major component of the hardened AS. It is responsible for the management of all the other components within the same hardened AS, the attack detection and traceback, and also the coordination of attack recovery with controllers of other ASes.

The Hardened AS Controller consists of a Sub-CA, a Configuration Base, a Key Base, a Manager Process, a Key Exchange Agent, a Rule Base, a Monitor Process, and a Recovery Process, shown in Fig 4. The Sub-CA issues certificates

to local routers. The Configuration Base stores the router configuration information. The Key Base saves the keys for decryption. The Manager Process manages the Key Base. The Key Exchange Agent is responsible for exchanging keys with hardened routers. The rule base stores rule information for detection. The Monitor Process monitors the traffic. The Recovery Process traces the source of the attack.

3.3. HARDENED NETWORK ROUTING --HBGP

As stated earlier, *the knowledge of the "last hardened router" is actually a very critical part of this hardening concept*. Since if we cannot figure out the last hardened router correctly, the packets will end up unrecognized and dropped.

All the border routers in an AS get up-to-date routing information by running Border Gateway Protocol (BGP [2]). We decided to *compatibly extend* BGP so that the hardened border routers would not only get the routing information from BGP messages, but also the last hardened AS information of each reachable route. So we can acquire this important information without significantly increasing the cost. We call this extension Hardened Border Gateway Protocol (HBGP).

In BGP, the routing information is transmitted using UPDATE [2] messages. In every UPDATE[2] message a variable length of sequences of Path Attributes is presented. In HBGP, we defined a new path attribute, HASPATH. In order to be compatible with BGP, this new attribute is also a triple, <attribute type, attribute length, attribute value>.

For HASPATH, the attribute type is set to be *optional, transitive, and partial*. Transitive and partial flags make *normal border routers propagate this attribute even though they do not understand it.* The attribute value always includes the last hardened AS information.

If a HBGP speaker advertises the route to its own autonomous system, it does not modify the HASPATH attribute. If advertises the route to a neighboring autonomous system, it needs to update the HASPATH as follows:

1. If no HASPATH is present, the router needs to attach the HASPATH attribute to the message with the local hardened AS controller IP address as the attribute value.
2. If the HASPATH is already present, do nothing.

We only use BGP for the purpose of Hardened Network information, not intent to secure BGP. Also the security of HBGP relies on that of BGP.

3.4. THE "FIRST/LAST HARDENED ROUTER"

For each new route, a hardened border router establishes a shared secret with the last hardened AS Controller of this route. If it does not have a valid key, its Key Exchange Agent contacts with the last hardened AS Controller to agree on a shared secret using SSL. Then the hardened border router stores the last hardened AS information in Local RIB and saves the shared key.

For every packet, the first hardened router the packet encounters encrypts the payload. But the route might change when network situation changes. If it happens before the new route is propagated and after the encryption has been done, the packet will be dropped due to no decryption key. In order to avoid it, the IP packet format sent out by the first hardened router is shown in Fig. 5.

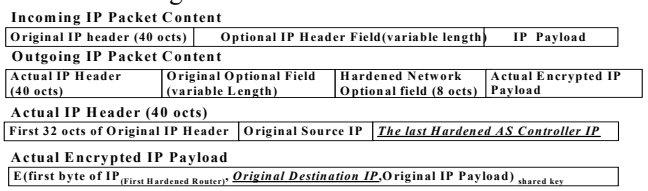

Figure 5. Actual IP Packet Content

The other hardened routers the packet passes through simply forward it. Only before the hardened router in the last Hardened AS on the route sends the packet to the neighboring AS, the router decrypts the packet, and matches the IP addresses in the payload and the header. If match, it restores and forwards the original packet; otherwise, it drops the packet.

Hardened access routers share keys with each other. They encrypt/decrypt the traffic only transmitted inside the AS.

Except these operations, the first/last hardened routers have to calculate the number of packets for each destination. They report them to the local AS Controller for flow control.

3.5. IMPLEMENTATION & SIMULATION

We implemented HBGP by extending the BGP implementation of Gated. The hardened network information gets propagated on the mixed hardened/normal Linux routers.

We implemented hardened router with 128-bit RC4 encryption algorithm. We set up a 100Mbps network using 3 1Ghz PCs. One PC worked as a normal/hardened router by changing to different kernels. One PC sent out TCP blocks with different sizes to the third machine. We compared the transmit delay and the router CPU usage, shown in Fig. 6 and 7.

In Fig. 6, there was not much difference between normal and hardened network. Because the router can encrypt in line speed by devoting more CPU resource, shown in Fig. 7. Hence the encryption is not the bottleneck.

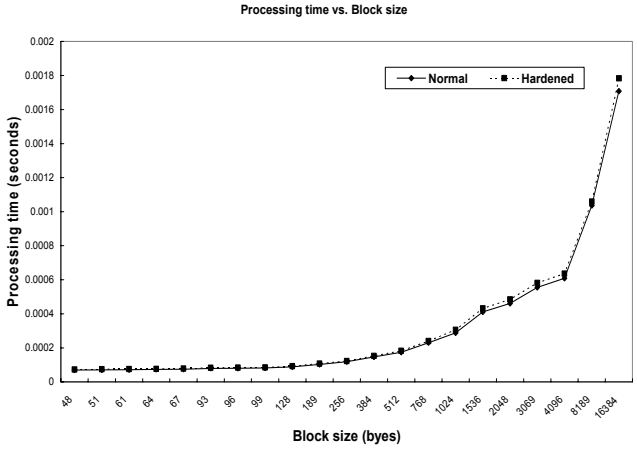

Figure. 6. Processing Delay of Normal and Hardened Network

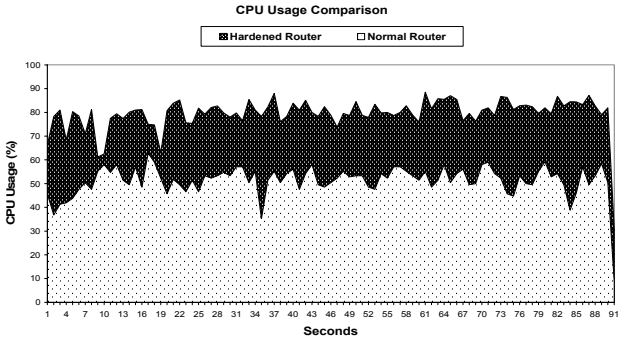

Figure. 7. CPU Usage Comparison of Normal/Hardened Network

We simulated the Hardened Network using SSFnet. We simulated a 100Mbps network with 26 ASes on a 1Ghz PC. Each AS had one router and one host. We measured the packet transmission delay by hardening any pair of backbone routers or end routers. Fig. 8 shows the comparison of these measurements, classified by the number of hops. Hardening the backbone routers introduced more overheads. Because Backbone routers have to handle more traffic; therefore more traffic will be affected.

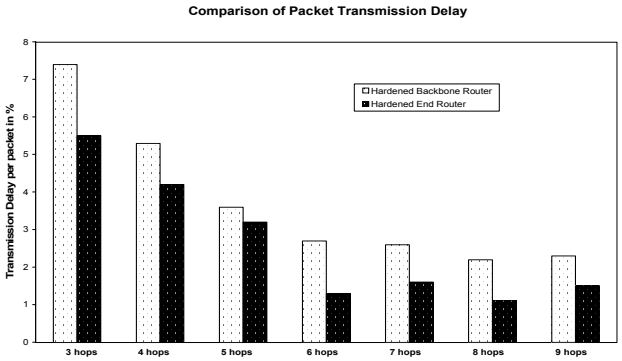

Figure. 8. Trans. Delay for hardening Backbone/end Routers

These graphs described the overhead the Hardened Network will introduce. It is caused by encryption/decryption and the traffic statistics, which are the techniques we use to detect and trace the source of DDoS. The network will benefit from these functionalities when more and more routers are hardened. We simulated a Ping attack to one of the ASes, the victim to see how the Hardened Network responded.

Since the detail algorithms of information gathering and traffic control are illustrated in [9], here we only use Fig. 9 to demonstrate the benefit of the Hardened Network.

Fig. 9 shows the traffic that the victim got under normal and attack situations for hardening all the 18 end routers and only hardening the gateway of the victim. It indicates that with the traffic statistics, we can detect the DDoS attack at the earlier stage, and also we can notify the first hardened router to drop the attack traffic using the identification.

Also the performance of the Hardened Network depends significantly on the speed of encryption. Currently, we are working to build routers using the Intel IXP 1200 network processor [8], but the work has not been completed.

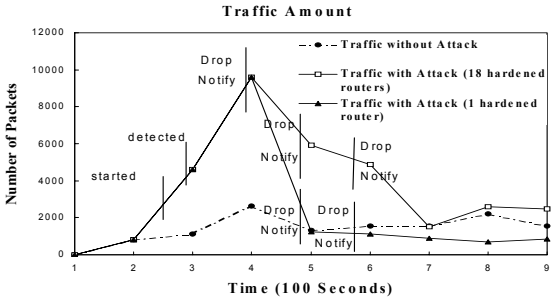

Figure. 9. Traffic Amount at Server end

4. CONCLUSION

In this paper, we discussed a new approach that can be implemented to provide infrastructure-based security to a wide area, large-scale open networking, such as the Internet. The methods outlined can be incrementally deployed and are effective against a wide range of attacks, including DDoS attack. The approach calls for embedding cryptographic and control protocols into the fabric of the network, and keeping a consistent, compatible TCP/IP interface to all un-upgraded networking ASes. Our preliminary implementation and performance studies show the feasibility of this approach.

REFERENCES

[1] Steven M. Bellovin, "Security Problems in the TCP/IP Protocol Suite," AT&T Bell Laboratories, Murray Hill, New Jersey, April 1989.

[2] Y. Rekhter, "A Border Gateway Protocol 4," Internet Draft, RFC 1771, March 1995, Network Working Group.

[3] P. Ferguson, and D. Senie, "Network Ingress Filtering: Defeating Denial of Service Attacks which Employ IP Source Address Spoofing," RFC2827, May, 2000, Network Working Group.

[4] Michael Lyu, Lorrien Lau, "Firewall Security: Policies, Testing and Performance Evaluation," COMPSAC2000, IEEE 0-7695-0792-1/00.

[5] D. Maughan, M. Schertler, M. Schneider, and J. Turner. "Internet Security Association and Key Management Protocol," Internet Draft, draft-ietf-ipsec-isakmp-09.txt, March 1998.

[6] A. Frier, P. Karlton, and P. Kocher, "The SSL 3.0 Protocol," Netscape Communications Corp. November 18, 1996.

[7] Y. Ylonen, "SSH Transport Layer Protocol," Internet Draft, March 20,2002, Network Working Group.

[8] Intel Corporation, "IXP 1200 Network Processor Datasheet," Spetember 2000.

[9] Shu Zhang, Partha Dasgupta, "Denying Denial-of-Service Attacks: A Router Based Solution," appearing in the proceeding of 2003 International Conference on Internet Computing.

[10] CERT Coordination Center, "Denial of Service Attacks,"http://www.cert.org/tech_tips/denial_of_service.html.

[11] Dawn Xiaodong Song and Adrian Perring, "Advanced and Authenticated Marking Schemes for IP Traceback," in Proceedings of the 2001 IEEE Infocom Conference.

[12] S.M.Bellovin, ICMP Traceback Messages. Internet Draft: Draft-bellovin-itrace-00.txt, Mar. 2000.

[13] H.Burch, and B. Cheswick, "Traceing Anonymous Packets to Their Approximate Source", in proceeding of Usenix LISA'00